

REDUCTION OF CHROMATIC BLEEDING ARTIFACTS IN IMAGES CONTAINING SUBSAMPLED CHROMINANCE VALUES

5

BACKGROUND OF THE INVENTION

The present invention relates to digital images. More specifically, the present invention relates to the reconstruction of compressed images.

Image compression according to JPEG and certain other image compression algorithms is performed by converting images to luminance and chrominance channels and subsampling the chrominance channels (Cb and Cr) by an integer factor (usually 2 or 4). The luminance and subsampled chrominance channels are then compressed by a compression module suitable for grayscale image. The chrominance channels may be subsampled because the human visual system (HVS) has a reduced spatial sensitivity of chromatic sensation relative to luminance sensation by roughly a factor of two. Consequently, chromatic bleeding artifacts resulting from the sub-sampling are virtually unnoticeable in most color photographs that contain subsampled chrominance values.

However, chromatic bleeding artifacts are noticeable in color compound document images, which frequently contain sharp edges between different colors. For example, a color compound document might have text or other computer-generated shapes that are overlaid on regions having uniform color. The sharp edges occur at the transitions of the computer-generated images and the photographs.

The chromatic bleeding artifacts are noticeable if the reconstructed chrominance-channels are interpolated back to their original resolution by a simple technique such as linear interpolation. The chromatic bleeding artifacts are seen at boundaries between regions with saturated colors and white regions (or regions with bright colors). In document images, these boundaries are often between a white or bright background and color characters, lines and patches.

The chromatic bleeding artifacts are also noticeable in regions containing inverted text (e.g., white characters on a colored background). Due to the properties of the human visual system, the chromatic bleeding artifacts are visible on the brighter side of the edge.

5 The chromatic bleeding artifacts reduce image quality. Therefore, it would be desirable to reduce the chromatic bleeding artifacts, especially in color compound document images.

10 SUMMARY OF THE INVENTION

According to one aspect of the present invention, chromatic bleeding artifacts in a digital image are reduced by modifying chrominance values of at least some pixels in the digital image. The pixels are modified according to luminance values and chromatic dynamic range.

15 Other aspects and advantages of the present invention will become apparent from the following detailed description, taken in conjunction with the accompanying drawings, illustrating by way of example the principles of the present invention.

20

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is an illustration of a method of reconstructing an image and reducing chromatic bleeding artifacts in the reconstructed image;

Figure 2 is an illustration of a local neighborhood of pixels;

25 Figures 3(a) to 3(d) are illustrations of a sliding window for determining chromatic dynamic ranges of pixels in a reconstructed image that has been interpolated by pixel replication; and

Figure 4 is an illustration of a hardware implementation of the method of Figure 1.

30

DETAILED DESCRIPTION OF THE INVENTION

As shown in the drawings for purposes of illustration, the present invention is embodied in a method and apparatus for reducing chromatic bleeding artifacts in digital images containing subsampled chrominance values. The method and apparatus are particularly suited for reducing chromatic bleeding artifacts in color compound document images containing both natural and computer-generated features. Natural features in photo regions tend to have mild luminance or chrominance transitions. Computer-generated features such as text regions tend to have sharp edges and sharp luminance or chrominance transitions. The method and apparatus reduce the chromatic bleeding artifacts that occur at the sharp edges.

Reference is made to Figure 1, which illustrates a method of reconstructing a compressed image and reducing chromatic bleeding artifacts in the reconstructed image. The compressed image includes a first stream of compressed luminance (Y) values and second and third streams of compressed chrominance (Cb and Cr) values. Prior to compression, the chrominance channels were either subsampled or low-pass filtered and then subsampled.

The image reconstruction begins by accessing the first, second and third streams of compressed values (block 102). The luminance values (Y) are reconstructed from the first stream (block 104) and the chrominance values (Cb and Cr) are reconstructed from the second and third streams (block 106).

The luminance and chrominance values will hereinafter be described with respect to normalized values. For the luminance values, $0 \leq Y \leq 1$. For the chrominance values, $-0.5 \leq Cb \leq 0.5$ and $-0.5 \leq Cr \leq -0.5$. The point $Cb=Cr=0$ corresponds to the gray-axis (no color).

Since the chrominance values (Cb and Cr) were subsampled during compression, the reconstructed chrominance channels are interpolated or upsampled to their original resolution to match the resolution of the luminance channel (block 108). A simple interpolation method is pixel replication, where

every low-resolution pixel is replaced by a 2×2 block of high-resolution pixels having the same chrominance values (C_b and C_r). A more sophisticated upsampling method involves applying low-pass filtering on the pixel-replicated chrominance channels for reducing aliasing artifacts in locations of high spatial frequency, such as edges.

Chromatic bleeding artifacts in the reconstructed digital image are reduced by modifying chroma values of at least some of the pixels in the digital image. The chroma value of a pixel is modified with respect to the pixel's luminance value and local chromatic dynamic range (block 110); and the chrominance values (C_b and C_r) are scaled according to the change in chroma value (block 112). Modifying the chroma value with respect to chromatic dynamic range is performed so as not to reduce chromaticity too much in regions with essentially uniform light colors.

The reconstructed luminance and chrominance channels may be converted back to RGB space at the original resolution. The reconstructed image may then be displayed (e.g., printed).

The local chromatic dynamic range of a pixel is determined with respect to its local pixel neighborhood. Figure 2 provides an example of a neighborhood of pixels. The neighborhood, delineated by a window indicated in dashed lines, includes a 3×3 array of pixels. The pixel of interest, marked by an "X," lies at the center of the neighborhood.

The chromatic dynamic range for each pixel of interest may be computed as the difference between minimum and maximum chroma values of the pixels in the local neighborhood. The chroma value (C) of a pixel may be computed as $C = \sqrt{C_b^2 + C_r^2}$ or the chroma value (C) may be approximated by $C \approx |C_b| + |C_r|$ where the allowed value range for each of the chrominance values is $[-0.5, 0.5]$. The chromatic dynamic range of a neighborhood may be computed as $D = C_M - C_m$ where C_m represents the minimum chroma value of the neighborhood and C_M represents the maximum chroma value of the neighborhood.

A neighborhood and chromatic dynamic range are determined for each pixel of interest. Consequently, the chromatic dynamic range is determined on a pixel-by-pixel basis.

- A new chroma value (C') for a pixel of interest may be computed as
- 5 $C' = C - f(Y, D) \cdot (C - C_0)$, where C is the original chroma value for the pixel, D is the chromatic dynamic range, C_0 is a chromatic modulus having a value between zero and the minimum chroma of the local neighborhood for the pixel, and $f(Y, D)$ is a parametric expression that determines the amount of relative chroma reduction.

10 The parametric expression $f(Y, D)$ may range between 0 and 1. The parametric expression $f(Y, D)$ complies with the following constraints: (1) $f(Y, D) \rightarrow 1$ for $Y \rightarrow 1$, whereby $C' \rightarrow C_0$; and (2) $f(Y, D) \rightarrow 0$ for $D \rightarrow 0$. If $D \rightarrow 0$, then $C' \rightarrow C$, whereby the chroma value for the pixel is not modified.

15 The parametric expression $f(Y, D)$ may also comply with another constraint: $f(Y, D) \rightarrow 0$ for $D \rightarrow 0$ and $Y \rightarrow 1$. Thus, the chromatic dynamic range takes precedence.

If the chroma value is modified, each corresponding chrominance value (Cb and Cr) is scaled by the ratio C'/C . Thus, $Cb' = \frac{C'}{C} Cb$ and $Cr' = \frac{C'}{C} Cr$. This keeps the hue of the color unchanged.

20 Consider different cases for the chromatic modulus (C_0). In a first case $C_0 = C_m$, whereby $C' = C - f(Y, D) \cdot (C - C_m)$ and the new chroma value (C') lies between the original chroma value (C) and the neighborhood minimum (C_m). For high luminance $Y \rightarrow 1$, the new chroma value (C') approaches the local minimum (C_m). Thus $C' \rightarrow C_m$ as $Y \rightarrow 1$.

25 In a second case, $C_0 = \max(C_m - D, 0)$. Such a modulus C_0 allows the new chroma value (C') to fall below the local minimum (C_m). However, the chroma value is still dependent upon the chromatic dynamic range (D). The chromatic modulus (C_0) is not reduced below $C_m - D$. If $C_m < D$, the chromatic modulus (C_0) becomes zero. For chromatic dynamic ranges that are not

small, the situation $C_m < D$ corresponds to edges between “non-colored” and “saturated color” regions. In this situation, the chroma value may be reduced to zero if the “non-colored” side is close to white. On the other hand, the situation $C_m > D$ could correspond to an edge between a “mild-color” and a 5 “saturated-color” where the visibility of color-bleeding artifacts is not high to start with. In this situation, allowing the chroma value to be reduced to zero would reduce too much color in the lighter side of the edge. Putting a lower bound $C_m - D$ on the modified chroma value ensures that in color gradient regions (“mild edges”) where the chromatic dynamic range (D) is neither very 10 small or very large, the chroma value will not be reduced significantly when the luminance $Y \rightarrow 1$.

In the third case, $C_0=0$. For high luminance $Y \rightarrow 1$, the new chroma value (C') approaches 0. Thus $C' \rightarrow 0$ as $Y \rightarrow 1$.

The parametric expression may be expressed in different ways. An 15 exemplary parametric expression is $f(Y, D) = \max\left[1 - \alpha\left(\frac{1-Y}{D}\right), 0\right]$. The exemplary parametric expression always complies with the two constraints: $f(Y \rightarrow 1, D) \rightarrow 1$ and $f(Y, D \rightarrow 0) \rightarrow 0$, regardless of the value of the term α (where α is positive).

For a given value of D , there is a luminance value (Y) below which the 20 chromatic bleeding artifacts are not visible. If a chromatic bleeding artifact is not visible for a given D , the value of the term α may be chosen such that $f(Y, D) = 0$. For example, the chroma value is not modified if $Y < (1-D/\alpha)$. That is, $C' = C$ if $Y < T_Y(D)$, where threshold $T_Y(D) = (1-D/\alpha)$.

The term α may be a function of the chromatic dynamic range (D). A 25 term α may be measured for each value of D . Each term α may be measured from subjective tests (e.g., observing visibility of artifacts in patches).

In the alternative, the term α may be constant for all D . The value for the term α can be determined by trial and error. For example, $0.5 \leq \alpha \leq 1$ has been found to work well.

Thus disclosed is a method for reducing the chrominance magnitude in pixels that would generate a visual perception of chromatic bleeding. By taking luminance and local chrominance values into account, the method does not modify patches that have high luminance values and non-negligible chrominance values. A flat region might have relatively high luminance but also non-negligible chrominance. The method would not map such a region to white. A document image might have text boxes with light colored backgrounds. The method would not map the light backgrounds to white.

The method is particularly useful for reducing chromatic bleeding artifacts in compound color documents. However, the method is not so limited.

Although Figure 2 shows a 3x3 square-shaped neighborhood, the method is not limited to such a neighborhood. The neighborhood is not limited to any particular size. The number of pixels is not limited to nine. Although a fixed number of pixels in the neighborhood is preferred for all pixels of interest, the size of the neighborhood may be changed dynamically to accommodate a particular class of image region (e.g., text, graphics, natural features).

The neighborhood is not limited to any particular geometry, although square windows are preferred for performance reasons. For example, the shape of the neighborhood may be diamond-shaped. The minimum/maximum calculations may include the center pixel (that is, a full neighborhood may be used) or the minimum/maximum calculations may exclude the center pixel of interest (that is, a hollow neighborhood may be used).

The method has been described in connection with decomposing a digital image into visually significant channels (e.g., luminance and blue and red chrominance channels). However, the method is not limited to YCbCr color space. For example, the method may be applied in Yuv, Luv or Lab color space (where L represents luminosity).

In the method described above, much of the computational effort is devoted towards computing the chromatic dynamic range for each pixel in a neighborhood. However, the computational effort can be reduced, and the speed of the method increased, by using subsampled chrominance values

and thus taking advantage of redundancy.

Reference is now made to Figures 3(a) to 3(d), which show a sliding window for determining chromatic dynamic ranges of pixels in a reconstructed image that has been interpolated by pixel replication. During the interpolation,

5 a first pixel in the reconstructed luminance channel is replaced by a 2x2 block of pixels C1, a second pixel in the reconstructed luminance channel is replaced by a 2x2 block of pixels C2, a third pixel in the reconstructed luminance channel is replaced by a 2x2 block of pixels C3, and a fourth pixel in the reconstructed luminance channel is replaced by a 2x2 block of pixels

10 C4.

The pixels in every 2x2 block in the upsampled chrominance channels contain exactly the same chrominance value. As Figures 3(a) to 3(d) show, each 4x4 block in the upsampled chrominance channel which is aligned with 2x2 block boundaries, has four possibilities of accommodating 3x3 neighborhoods. For each of these four possibilities, the neighborhood minimum (or maximum), and the chromatic dynamic range of the neighborhoods are the same, since they are calculated from the same four chromatic modulus values. The four chromatic modulus values correspond to the four double-sized pixels, where each such pixel corresponds to a pixel in

15 the sub-sampled chrominance channels. Thus, the computational complexity is reduced to finding the maximum and minimum in 2x2 neighborhoods in the subsampled chrominance channels. This reduces by a factor of four the number of chrominance modulus calculations, and reduces by a factor of sixteen the number of comparisons needed for calculating the minimum and

20 maximum chroma values (C_m and C_M) of the neighborhoods. If pixel-replication is to be used in this manner, it is preferable to apply low-pass filtering on the chrominance channels prior to their subsampling. For example, every 2x2 block of the original chrominance channels could be replaced by the average or the median of the four chrominance-values in that block.

25 The computational effort may be reduced in other ways. For example, if the difference between the new chroma value (C') and the original chroma

value (C) is imperceptible, little will be gained by modifying the chroma value. Therefore, the parametric expression is not computed and the chroma value is not modified if the difference is imperceptible. For example, $C'=C$ if $(C-C_0) < T$, where T is a predetermined threshold. If the chromatic dynamic range (D) is 5 computed with a fast recursive algorithm, it is calculated for all pixels. Therefore, little overhead is incurred by performing the comparison $D < T$. An exemplary threshold is $T=0.1$ (for normalized chrominance values).

The method may be implemented in hardware, software or a combination of the two.

10 Reference is now made to Figure 4, which illustrates a software implementation of the method. Computer memory 212 stores a program 214 that, when executed, instructs a processor 216 to perform the method described above. The memory 212 and the processor 216 may be part of a personal computer or workstation.

15 A non-recursive algorithm may be used to determine the minimum and maximum chroma values of pixel neighborhoods. The non-recursive algorithm may be used for all types of neighborhoods. The size of the memory needed to determine the minimum and maximum values is equal to the size of the neighborhood.

20 A fast recursive algorithm may be used to determine the minimum and maximum chroma values of full square neighborhoods (e.g., 2x2 neighborhoods, 3x3 neighborhoods, 5x5 neighborhoods). Except for the 2x2 neighborhood, only twelve comparisons are needed to determine the maximum and minimum chroma values of a neighborhood, regardless of the 25 size of the neighborhood. In the case of the 2x2 neighborhood, only three comparisons are needed. If this recursive algorithm is used, the size of memory needed to buffer the lines is equal to twice the size of the lateral kernel window.

The present invention is not limited to the specific embodiments 30 described and illustrated above. Instead, the invention is construed according to the claims that follow.